



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/726,948	12/02/2003	Stephen C. Hahn	SUNMP453	5353
32291	7590	10/02/2007	EXAMINER	
MARTINE PENILLA & GENCARELLA, LLP			WAI, ERIC CHARLES	
710 LAKEWAY DRIVE			ART UNIT	PAPER NUMBER
SUITE 200			2195	
SUNNYVALE, CA 94085				
MAIL DATE		DELIVERY MODE		
10/02/2007		PAPER		

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Office Action Summary	Application No.	Applicant(s)	
	10/726,948	HAHN ET AL.	
	Examiner	Art Unit	
	Eric C. Wai	2195	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) Responsive to communication(s) filed on 02 December 2003.
- 2a) This action is FINAL. 2b) This action is non-final.
- 3) Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) Claim(s) 1-53 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) Claim(s) _____ is/are allowed.
- 6) Claim(s) 1-53 is/are rejected.
- 7) Claim(s) _____ is/are objected to.
- 8) Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) The specification is objected to by the Examiner.
- 10) The drawing(s) filed on 02 December 2003 is/are: a) accepted or b) objected to by the Examiner. Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a). Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) All b) Some * c) None of:
 1. Certified copies of the priority documents have been received.
 2. Certified copies of the priority documents have been received in Application No. _____.
 3. Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- 1) Notice of References Cited (PTO-892)
- 2) Notice of Draftsperson's Patent Drawing Review (PTO-948)
- 3) Information Disclosure Statement(s) (PTO/SB/08)
Paper No(s)/Mail Date 6/20/2005
- 4) Interview Summary (PTO-413)
Paper No(s)/Mail Date. _____
- 5) Notice of Informal Patent Application
- 6) Other: _____

DETAILED ACTION

1. Claims 1-53 are presented for examination.

Claim Rejections - 35 USC § 112

2. The following is a quotation of the second paragraph of 35 U.S.C. 112:

The specification shall conclude with one or more claims particularly pointing out and distinctly claiming the subject matter which the applicant regards as his invention.

3. Claims 1-31, and 42-53 are rejected under 35 U.S.C. 112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which applicant regards as the invention.

- a. The following terms lack antecedent basis in the claims:

- i. Claim 16, line 2, "said workload".

- b. The following claim terms are not clearly understood:

- ii. Claims 1, 42 and 48 recite, "determining a range of computer resources to make available for other use". It is unclear what is meant by other use (e.g. other use by the workload, or other use by other workloads in the system).

- iii. Claim 9 recites, "wherein said monitoring is performed by a process within a user space". It is unclear whether a process within the user space is performing the monitoring or whether a process only monitors the user space.

iv. Claim 16 recites, "wherein said selecting occurs within said user space". It is unclear whether a process within the user space is performing the selecting or whether a process only selects from within the user space.

v. Claim 26 recites, "wherein said monitoring and said determining occur within said workload". It is unclear whether a process within the user space is performing the monitoring and determining or whether a process only monitors and determines within the user space.

Claim Rejections - 35 USC § 103

4. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

5. Claims 1-35, and 38-53 are rejected under 35 U.S.C. 103(a) as being unpatentable over Applicant's Admitted Prior Art (AAPA) in view of Tanenbaum (Modern Operating Systems, 2nd edition, 2001, Prentice Hall Intl, New Jersey)

6. Tanenbaum was disclosed in IDS dated 06/20/2005.

7. Regarding claim 1, AAPA teaches a computer implemented method of managing computer resources comprising:

accessing an amount of computer resources allocated to a workload of a computer system (pg 3 lines 14-16);

monitoring computer resource usage of said workload (pg 3 lines 14-16, wherein the request is determined to request more physical memory than is available); and

determining a range of computer resources to make available for other use provided said computer resource usage of said workload exceeds said amount of computer resources allocated to said workload (pg 3 lines 16-20, wherein in response to the overcommit, the system will free up resources for the requestor).

8. AAPA does not explicitly teach wherein said monitoring and said determining occur within a user space. In fact, AAPA teaches the steps of monitoring and determining for all resources (i.e. all user spaces).

9. Tanenbaum teaches the use of local versus global allocation policies. (pg 234, "4.6.1 Local versus Global Allocation Policies"). By using a local allocation policy, the system does not need to continually decide how many page frames to assign each process.

10. It would have been obvious to one of ordinary skill in the art at the time of the invention to perform the monitoring and determining steps within a user space. Since the working set size may change in size quickly, as indicated by Tanenbaum, one would be motivated by the desire to use a strict allocation policy to reduce the need to continually decide how many pages to assign each process.

11. Regarding claim 2, AAPA teaches that the computer resources comprise physical memory (pg 3 lines 14-15).
12. Regarding claim 3, AAPA teaches that said determining comprises determining if a page of physical memory utilized by said workload has been accessed by said workload within a predetermined period of time (pg 4 lines 20-22).
13. Regarding claim 4, AAPA and Tanenbaum do not explicitly teach determining if a page of physical memory utilized by said workload has been accessed by said workload within a predetermined period of time comprises determining if said page has been accessed by said workload since a previous determination of whether said page had been accessed by said workload.
14. It would have been obvious to one of ordinary skill in the art at the time of the invention, to repeatedly make the determination whether a page in the physical memory has been recently accessed. One would be motivated by the desire to continuously monitor the usage of each page within the physical memory.
15. Regarding claim 5, AAPA teaches that computer resources comprises virtual memory (pg 3 lines 23-24).

16. Regarding claims 6-8, AAPA and Tanenbaum do not teach that the computer resources comprises central processing unit time, input/output space, or network bandwidth.

17. However, it would have been obvious to one of ordinary skill in the art at the time of the invention to include computer resources comprising central processing unit time, input/output space, or network bandwidth. It is well known in the art that computers are constrained by such resources.

18. Regarding claim 9, AAPA teaches a computer implemented method comprising: monitoring usage of a computing resource utilized by a workload (pg 3 lines 14-16, wherein the request is determined to request more physical memory than is available);

responsive to exceeding a limit on utilization of said computing resource, decreasing usage of said computing resource by said workload (pg 5 lines 1-9, wherein a given workload is victimized when it exceeds its allocation).

19. AAPA does not explicitly teach wherein said monitoring occurs within a user space. In fact, AAPA teaches the step of monitoring for all resources (i.e. all user spaces).

20. Tanenbaum teaches the use of local versus global allocation policies. (pg 234, "4.6.1 Local versus Global Allocation Policies"). By using a local allocation policy, the system does not need to continually decide how many page frames to assign each process.

21. It would have been obvious to one of ordinary skill in the art at the time of the invention to perform the monitoring step within a user space. Since the working set size may change in size quickly, as indicated by Tanenbaum, one would be motivated by the desire to use a strict allocation policy to reduce the need to continually decide how many pages to assign each process.

22. Regarding claim 10, AAPA teaches that said computing resource comprises physical memory (pg 3 lines 14-15).

23. Regarding claim 11, AAPA teaches that decreasing usage of said computing resource comprises paging a portion of said physical memory assigned to said workload out of said physical memory (pg 4 lines 7-13).

24. Regarding claim 12, AAPA teaches that said portion of said physical memory comprises a least recently used portion of said physical memory assigned to said workload (pg 4 lines 20-22).

25. Regarding claim 13, AAPA teaches that said decreasing usage does not halt operation of said workload (pg 5 lines 8-9, wherein the workload will continue operation).

26. Regarding claim 14, AAPA teaches that said decreasing usage is initiated by a process of said workload (pg 5 lines 18-20).

27. Regarding claim 15, AAPA and Tanenbaum do not teach that said process that performs said monitoring is not an operating system kernel process.

28. AAPA does teach that operating system kernel processes can be very expensive in terms of computing load on a system as they execute very frequently (pg 5 lines 13-15). It would have been obvious to one of ordinary skill in the art to perform the monitoring by some other process. One would be motivated by the desire to reduce computing load.

29. Regarding claim 16, AAPA teaches a computer implemented method comprising: accessing a list of memory pages assigned to said workload (pg 5 lines 11-12); responsive to a request from a first process of said workload for memory which exceeds a predetermined memory limit for said workload, selecting a plurality of memory pages from said list of memory pages (pg 5 lines 11-12); and initiating a second process within a user space to page out said plurality of memory pages (pg 5 lines 11-13).

30. AAPA does not explicitly teach wherein said selecting occurs within a user space. In fact, AAPA teaches the step of selecting for all resources (i.e. all user spaces).

31. Tanenbaum teaches the use of local versus global allocation policies. (pg 234, "4.6.1 Local versus Global Allocation Policies"). By using a local allocation policy, the system does not need to continually decide how many page frames to assign each process.
32. It would have been obvious to one of ordinary skill in the art at the time of the invention to perform the selecting step within a user space. Since the working set size may change in size quickly, as indicated by Tanenbaum, one would be motivated by the desire to use a strict allocation policy to reduce the need to continually decide how many pages to assign each process.
33. Regarding claim 17, AAPA teaches that accessing, selecting and initiating are performed by said second process within said workload (pg 5 lines 11-20, wherein the steps are performed by another process).
34. Regarding claim 18-19, AAPA and Tanenbaum do not teach that said second process is not an operating system kernel process or that said second process is loaded into a user space.
35. AAPA does teach that operating system kernel processes can be very expensive in terms of computing load on a system as they execute very frequently (pg 5 lines 13-15). It would have been obvious to one of ordinary skill in the art to perform the monitoring by some other process. One would be motivated by the desire to reduce computing load.

36. Regarding claims 20-21, AAPA teaches that said plurality of memory pages comprises memory pages that are least recently used by said workload (pg 4 lines 20-23).

37. Regarding claim 22, AAPA and Tanenbaum do not explicitly teach that said page out of said plurality of least recently used memory pages reduces a number of memory pages assigned to said workload to below said memory limit.

38. However, it would have been obvious to one of ordinary skill in the art at the time of the invention, that if the page out was performed in response to the physical requirements being exceeded, that the course of action would remedy the excessive use of resources.

39. Regarding claim 23, AAPA and Tanenbaum do not teach that said plurality of least recently used memory pages comprises the minimum number of memory pages to reduce said number of memory pages assigned to said workload below said memory limit.

40. However, it would have been obvious to one of ordinary skill in the art at the time of the invention, that if the page out was performed in response to the physical requirements being exceeded, that the course of action would remedy the excessive use of resources.

41. Regarding claim 24, AAPA teaches wherein at least a portion of said workload continues to operate subsequent to said initiating (pg 5 lines 8-9, wherein the workload will continue operation).

42. Regarding claim 25, AAPA and Tanenbaum do not teach that said initiating is not performed by an operating system kernel process.

43. AAPA does teach that operating system kernel processes can be very expensive in terms of computing load on a system as they execute very frequently (pg 5 lines 13-15). It would have been obvious to one of ordinary skill in the art to perform the monitoring by some other process. One would be motivated by the desire to reduce computing load.

44. Regarding claims 26-31, they are the computer-usable medium claims of claims 1-4 above. Therefore, they are rejected for the same reasons as claims 1-4 above.

45. Regarding claim 32, AAPA teaches a computer implemented method comprising: scanning pages for a first workload of a computer to determine if each of said pages was accessed since a last scan (pg 4 lines 20-23, wherein the selecting checks all memory pages from throughout the entire computer system); repeating said scanning and said setting for a second workload (pg 4 lines 20-23, wherein the selecting checks all memory pages from throughout the entire computer system).

46. AAPA does not teach setting bits indicative of the result of said scanning within a scoreboard related to said first workload and wherein the scanning and setting occur within a user space of said computer.

47. Tanenbaum illustrates a table used to indicate the local page replacement algorithm (pg 235, Fig 4-28). It would have been obvious to one of ordinary skill in the art at the time of the invention to utilize some scoreboard or table to quickly assess which pages were the least recently used as indicated by Tanenbaum.

48. Furthermore, Tanenbaum teaches the use of local versus global allocation policies. (pg 234, "4.6.1 Local versus Global Allocation Policies"). By using a local allocation policy, the system does not need to continually decide how many page frames to assign each process.

49. It would have been obvious to one of ordinary skill in the art at the time of the invention to perform the scanning and setting step within a user space. Since the working set size may change in size quickly, as indicated by Tanenbaum, one would be motivated by the desire to use a strict allocation policy to reduce the need to continually decide how many pages to assign each process.

50. Regarding claim 33, AAPA and Tanenbaum do not teach scanning comprises checking said pages according to an order inherent to a list of said pages.

51. However, it is well known in the art to scan a list in order.

52. Regarding claim 34, AAPA teaches paging out a plurality of pages utilized by said first workload responsive to said determining (pg 4 lines 20-22).

53. Regarding claim 35, AAPA teaches determining if the number of pages utilized by said workload exceeds a predetermined limit (pg 4 lines 7-10).

54. Regarding claim 38, AAPA teaches a computer implemented method of managing computer resources over a plurality of workloads, said method comprising:
for each workload of said plurality of workloads; monitoring respective workload resource usage against a respective allotment of each workload (pg 5 lines 11-12);
determining a range of computer resources to page out for each workload whose resource usage exceeds its respective allotment (pg 5 lines 11-12); and
initiating a paging out operation of said range of computer resources and wherein said monitoring (pg 5 lines 11-13).

55. AAPA does not explicitly teach wherein said determining and initiating occur within a user space. In fact, AAPA teaches the step of selecting for all resources (i.e. all user spaces).

56. Tanenbaum teaches the use of local versus global allocation policies. (pg 234, "4.6.1 Local versus Global Allocation Policies"). By using a local allocation policy, the system does not need to continually decide how many page frames to assign each process.

57. It would have been obvious to one of ordinary skill in the art at the time of the invention to perform the determining and initiating steps within a user space. Since the working set size may change in size quickly, as indicated by Tanenbaum, one would be motivated by the desire to use a strict allocation policy to reduce the need to continually decide how many pages to assign each process.

58. Regarding claim 39, AAPA teaches that determining comprises determining least recently used pages for each workload whose resource usage exceeds its respective allotment (pg 5 lines 11-13).

59. Regarding claim 40, AAPA teaches that the process is situated within a workload of said plurality of workloads (pg 3 lines 22-24).

60. Regarding claim 41, AAPA teaches paging out said range of computer resources and wherein each workload whose resource usage exceeds its respective allotment remains partially operable during said paging out of its respective range of computer resources (pg 5 lines 1-9).

61. Regarding claim 42, AAPA and Tanenbaum substantially teach the claim according to the reasons given in claim 1 above. However, AAPA and Tanenbaum fail to teach a bus for functionally coupling elements of said computer system; physical

memory coupled to said bus for storing processor instructions and data; and a processor coupled to said bus.

62. However, official notice is made that it is well known in the art that a computer system consists of a processor, memory, and buses interconnecting elements.

63. Regarding claims 43-47, they are the computer system claims of claims 2, and 5-8 above. Therefore, they are rejected for the same reasons as claims 2, and 5-8.

64. Regarding claim 48, AAPA and Tanenbaum substantially teach the claim according to the reasons given in claim 1 above. However, AAPA and Tanenbaum fail to teach that the steps of accessing, monitoring, and determining are performed by a first computer on workloads belonging to a second computer.

65. AAPA teaches that such steps are performed by the operating system but that such operations are costly in terms of computing load (pg 5 lines 12-15). Therefore, it would have been obvious to one of ordinary skill in the art at the time of the invention to offload the monitoring steps to another computer. One would be motivated by the desire to reduce the computing load on a single system.

66. Regarding claims 49-53, they are the computer system claims of claims 2, and 5-8 above. Therefore, they are rejected for the same reasons as claims 2, and 5-8.

67. Claims 36-37 are rejected under 35 U.S.C. 103(a) as being unpatentable over Tanenbaum (Modern Operating Systems, 2nd edition, 2001, Prentice Hall Intl, New Jersey) in view of Applicant's Admitted Prior Art (AAPA).

68. Tanenbaum was disclosed in IDS dated 06/20/2005.

69. Regarding claim 36, Tanenbaum teaches a computer implemented method comprising:

accessing memory usage for a workload and examining page usage for each process of said workload (pg 234, "4.6.1 Local versus Global Allocation Policies", paragraphs 2-3, wherein the algorithm tries to find the least recently used page for all the processes);

aggregating usage of said each process to determine an aggregate usage for said workload (wherein it is inherent that a workload comprises many processes);

determining least recently used pages based on accessed bits associated with said workload (pg 235, Fig 4-28, wherein it is inherent that bits are set to indicate least recently used pages).

70. Tanenbaum does not teach performing the determining step based on whether the aggregate usage exceeds said memory utilization limit for said workload and supplying a range of least recently used pages in a system call to an operating system kernel for evicting said range of least recently used pages to reduce resource usage by said workload; and retaining at least partial operation of said workload during said page evicting.

71. AAPA teaches paging out memory in response to workloads exceeding their respective allotments whereby an operating system kernel process will evict selected pages (pg 5 lines 11-20).

72. It would have been obvious to one of ordinary skill in the art at the time of the invention to perform the determination in response to an aggregate usage exceeding the memory utilization limit. One would be motivated by the desire to ensure that overall memory usage does not exceed the amount that is physically available.

73. AAPA teaches that said determining and said supplying occur in a plurality of user space processes (pg 3 line 24 to pg 4 line 1, wherein multiple user spaces can exist in a system).

Conclusion

74. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Eric C. Wai whose telephone number is 571-270-1012. The examiner can normally be reached on Mon-Thurs, 9am-5pm.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Meng - Ai An can be reached on 571-272-3756. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

EW


MENG-ALT AN
SUPERVISORY PATENT EXAMINER
TECHNOLOGY CENTER 2100